



HP Reference Configurations for Oracle 10g Data Warehousing

Optimized for 5TB and 10TB raw data sizes

Standard components, predictable performance and rapid deployment for data warehouses between 5TB and 50TB expanded database size

Challenges

IT planners and architects have a wide range of technology options for their business intelligence (BI) applications. Each hardware platform decision requires them to research, and then balance, the tradeoffs between price, performance, long-term scalability, manageability, and vendor reputation. When rapid deployment is a key factor, the choice often spans dedicated BI hardware, rigidly prescribed single-vendor configurations, and building custom solutions from well-understood technologies. Such a range in choices often means complex, time-consuming, decision cycles that drain staff time to define and manage the benchmarks or proof-of-concepts. All of this takes away from the basic function of IT, which is to deliver information quickly to business users so that they can make sound, timely decisions.

Solution

To reduce the risk and deployment time for BI solutions, HP has created a comprehensive set of BI Reference Configurations. The HP BI Reference Configurations draw on extensive lab, customer benchmark, and industry standard benchmark investments as well as global experience supporting thousands of data warehouse and business intelligence customers. The configurations enable IT planners and architects to start from a “best-fit” reference configuration that reflects their raw data size, database, operating system, processor, and node scaling preference. The configuration can then be optimized based on their specific workload and requirements.

Each configuration reflects “real-world” BI customer environments and workloads, balancing processing power, storage, and throughput to provide cost-effective performance. HP BI Reference Configurations are modular and can easily scale up by adding CPUs, nodes, or storage on an as needed basis. The configurations have been defined in collaboration with Oracle so that customers can expect predictable high performance based on the chosen configuration.

The configurations are also designed to avoid bottlenecks and achieve balanced I/O throughout all system components, from server to storage, for optimal BI performance. Achieving balanced system I/O throughput specifically designed to satisfy the requirements of BI workloads was a key objective in developing the configurations because it maximizes system BI performance while optimizing ROI through proper utilization.

Some vendors encourage customers to purchase dedicated, proprietary hardware that cannot be repurposed as business needs change. Others recommend solutions from a narrowly prescribed set of single-vendor components. Neither approach delivers on the promise of open, flexible, IT solutions based on standardized technologies. Because HP strongly supports customer needs for flexibility and re-use options, HP BI Reference Configurations combine architectural choice with predictable performance and standards-based technology.

For convenience, these configurations are arranged into size groups, each scoped for a limited range of raw data sizes. Within each group, HP offers a rich set of architecture and performance options. This document focuses on the 5TB and 10TB raw data size configurations.

Table 1: Database size, user concurrence, and price

	250GB Raw Data Size	500GB Raw Data Size	1TB Raw Data Size	5TB Raw Data Size	10TB Raw Data Size
Database Instance Size	200GB — 1TB	500GB — 2TB	1TB — 5TB	5TB — 20TB	10TB — 50TB
Concurrent Users	30 — 80	40 — 120	60 — 180	100 — 300	150 — 400

The database instance size is commonly larger, in some cases as much as 5x larger, than the raw data size. This “expansion factor” can vary widely depending on customer requirements, aggregates, indexes, and summaries, causing the instance size to expand. The choice of deploying Oracle compression or partitioning can also have a large effect in reducing the instance size. In addition, space must be allocated for temporary tablespace and log files.

Each configuration defines a matched set of servers and storage, which balances I/O throughput seamlessly across all components to yield consistently high BI performance. I/O throughput is recognized as one of the best indicators of potential BI query performance. All are designed for Oracle Database 10g and also include Oracle Real Application Clusters (RAC) for clustered (scale-out) configurations. Each BI reference configuration group provides architecture options to meet a range of existing IT platform standards, coupled to storage products from the HP MSA, EVA and XP product families.

Within each group, configurations are provided for “standard workload” and “heavy workload” use cases. Both types are based on real-world observation of customer workloads. Standard workload configurations deliver performance suitable for most commonly observed workloads. Heavy workload

configurations are optimized for more intense BI workloads and higher user concurrency. Consult your HP performance specialist for advice on adjusting a configuration to meet specific needs.

For each group, HP offers practical configurations to support the most common customer architectures.

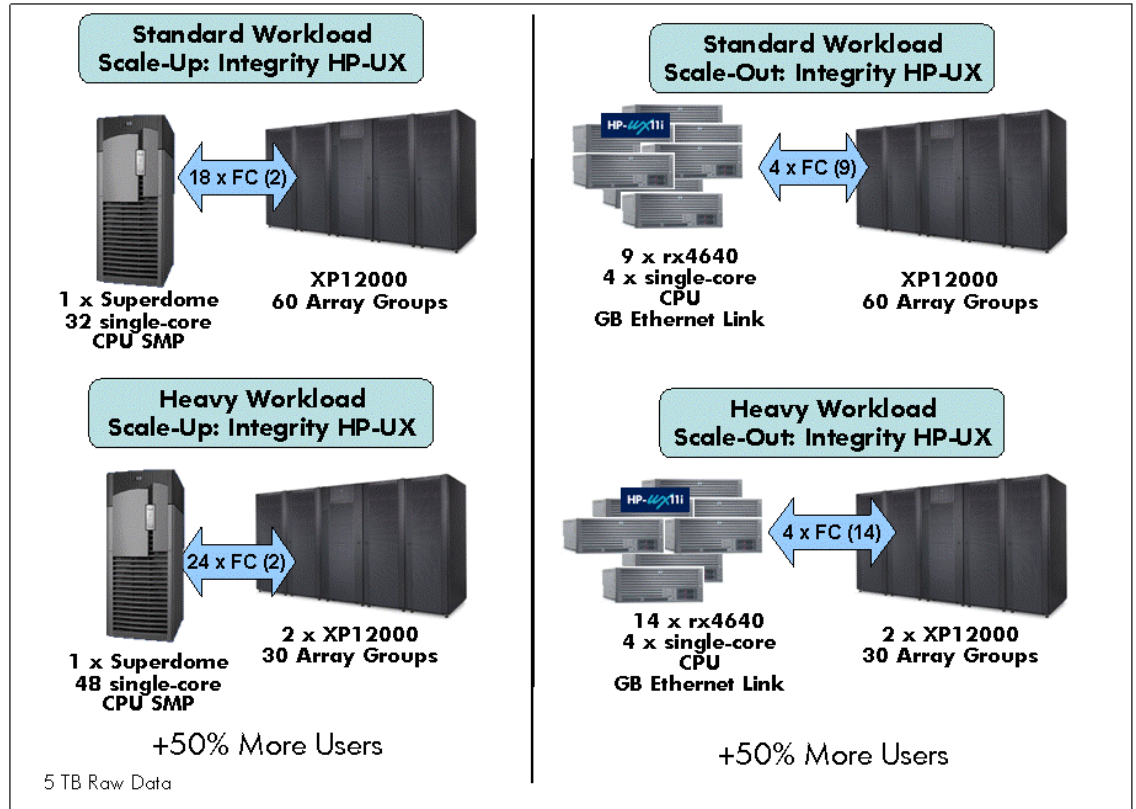
- **SMP / Scale-Up.** For customers who prefer the robust system management, flexibility, availability, and scaling provided by 64-bit HP Integrity systems and the HP-UX operating system (OS). SMP configurations can be run as dynamic partitions in large HP Integrity systems. Through the use of Instant Capacity on Demand (iCAP), customers can improve processor utilization and meet peak performance demands. Scale-up configurations support the Oracle Database but do not require Oracle RAC. High availability, partitioning, and other server and storage management are handled through the HP Virtual Server Environment (VSE), HP OpenView and HP Serviceguard products.
- **Cluster / Scale-Out with HP Integrity/HP-UX.** For customers who prefer to cluster small to mid-range 64-bit HP Integrity systems with the HP-UX OS. This configuration leverages the latest Oracle RAC technology for flexible node, CPU, and storage load and failover management.
- **Cluster / Scale-Out with HP Integrity/Linux.** For customers who wish to take advantage of the HP Integrity platform with the RH Linux OS. This configuration leverages the latest Oracle RAC technology for flexible node, CPU, and storage load and failover management.
- **Cluster / Scale-Out with HP ProLiant/Linux.** For customers who prefer to cluster multiple Opteron-based HP ProLiant systems with RH Linux. This configuration leverages the latest Oracle RAC technology for flexible node, CPU, and storage load and failover management.

In addition to taking into account the typical Oracle data expansion of tables, summaries, indexes, and views, the HP BI Reference Configurations for Oracle 10g also provide for RAID-based storage availability and data spreading across disks to maximize I/O throughput and performance. Storage array type and sizing were specifically chosen for best fit with the workload intensity, OS, server, and price considerations. See your HP server or storage performance specialist for more information.

HP 5TB HP-UX reference configurations for Oracle

The scale-out configurations employ Oracle RAC. XP12000 arrays were selected for best price/performance with HP-UX. The heavy workload configurations benefit from increased I/O bandwidth and data spreading enabled by multiple XP frames.

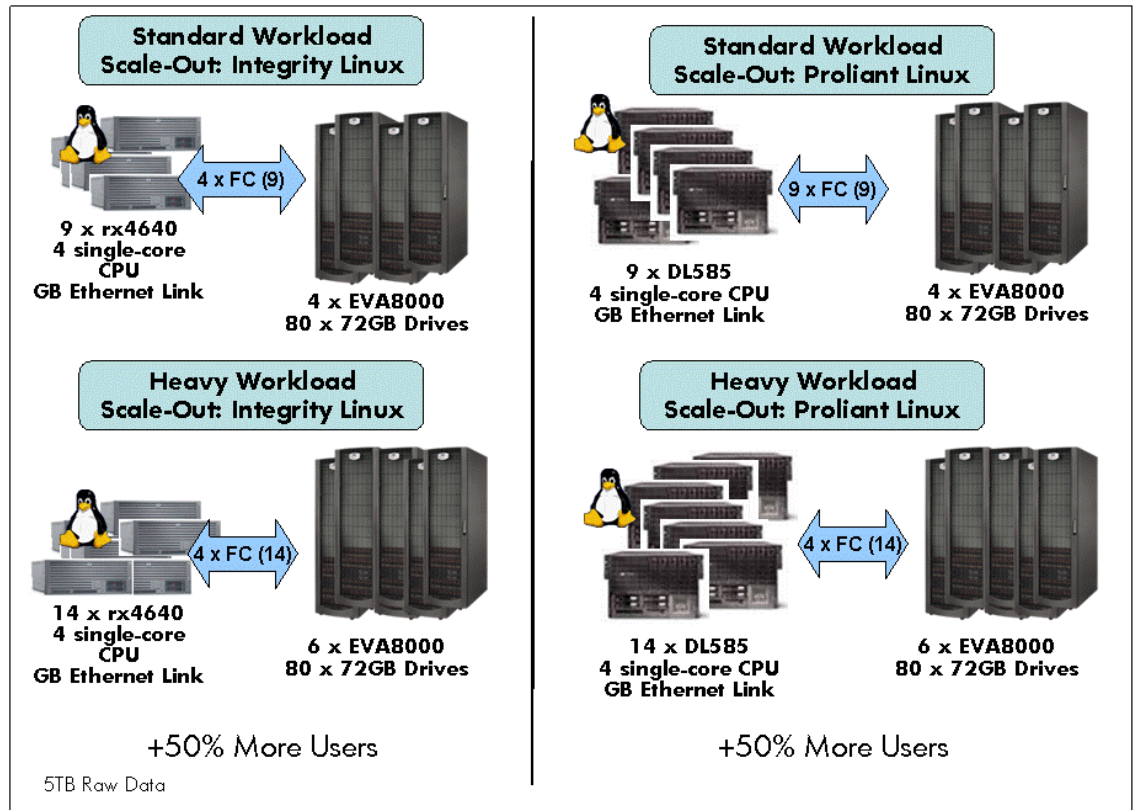
Figure 1: HP 5TB HP-UX reference configurations for Oracle



HP 5TB Linux reference configurations for Oracle

Oracle RAC is part of all configurations. Heavy workloads are addressed with larger clusters, plus additional EVA arrays for I/O bandwidth and more data spread. If customers choose, they can replace the EVA arrays in these configurations with the XP arrays from the HP-UX configurations.

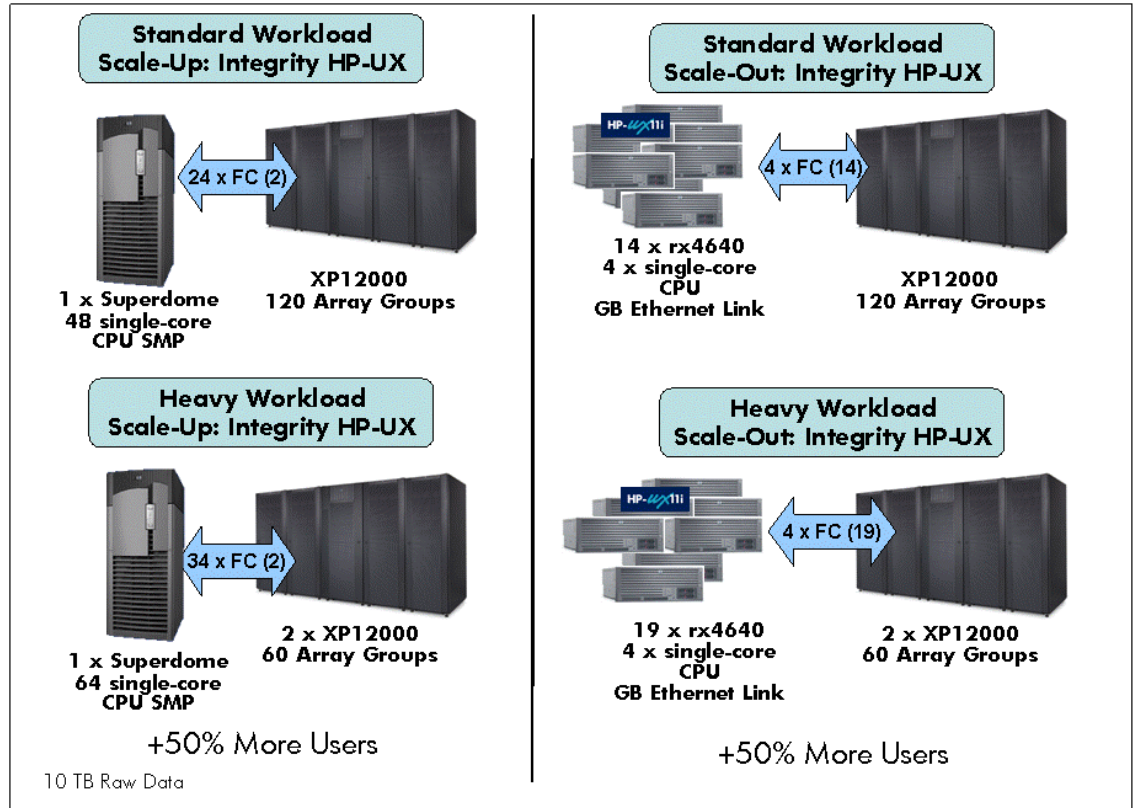
Figure 2: HP 5TB Linux reference configurations for Oracle



HP 10TB HP-UX reference configurations for Oracle

All configurations rely on XP12000 storage for I/O performance. Heavy workload configurations add nodes and / or CPUs and split the array groups across two XP12000 frames.

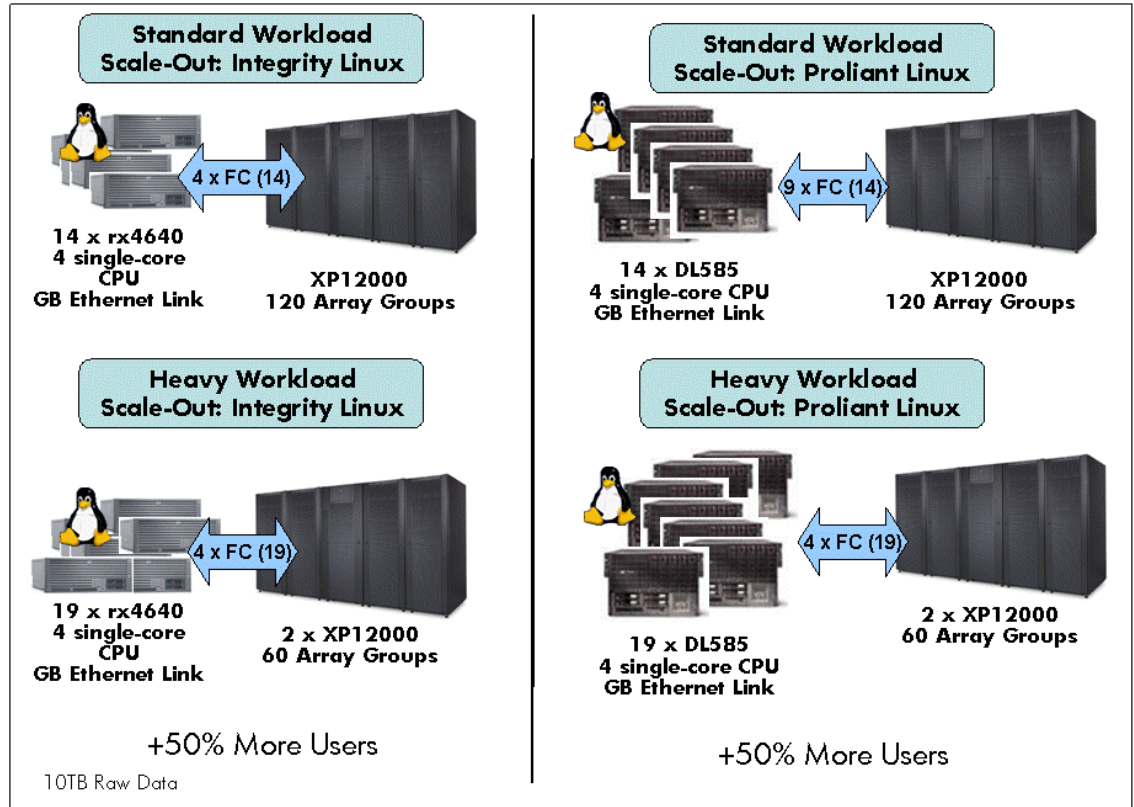
Figure 3: HP 10TB HP-UX reference configurations for Oracle



HP 10TB Linux reference configurations for Oracle

All configurations rely on XP12000 storage for I/O performance. Heavy workload configurations add nodes and / or CPUs and split the array groups across two XP12000 frames.

Figure 4: HP 10TB Linux reference configurations for Oracle



Pricing

The actual cost of the different architecture and OS options by performance group depends on customer pricing agreements with HP and Oracle. Relative price breakdowns for the 10TB groups are shown below.

Figure 5: Relative costs for 10TB raw data, standard workload performance group

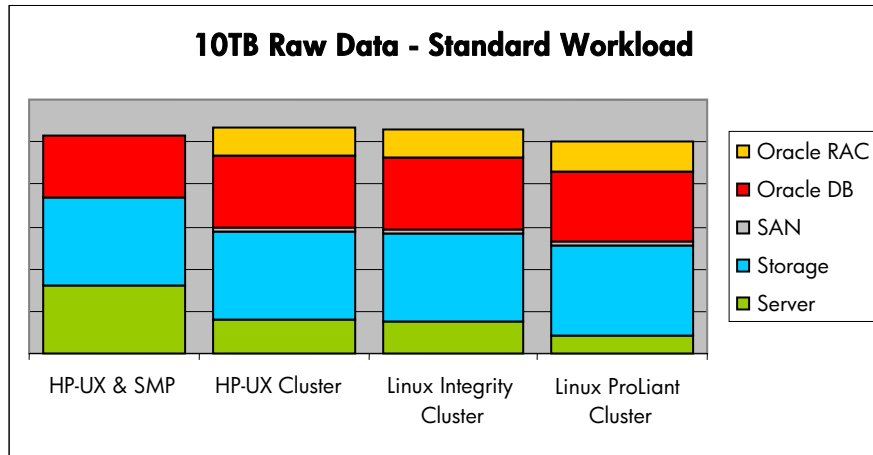
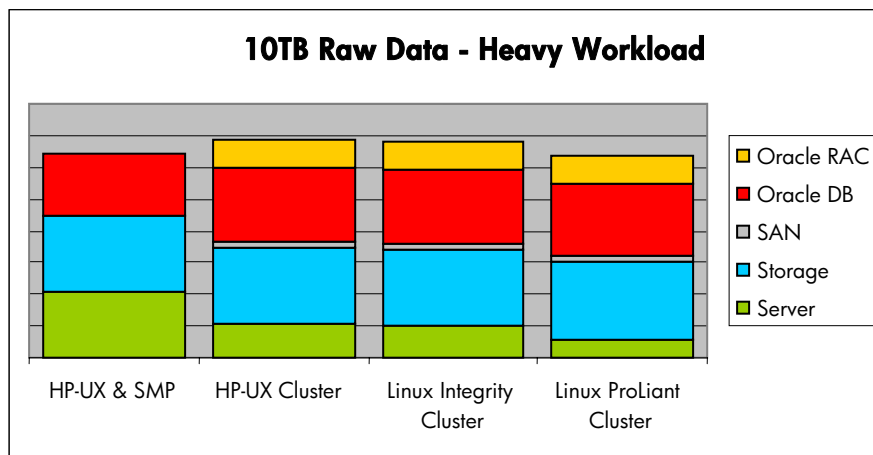


Figure 6: Relative costs for 10TB raw data, heavy workload performance group



Benefits

IT staff time and resources are precious. HP BI reference configurations enable IT departments to do more with less, more quickly, and at lower risk. Some immediate benefits for deploying data warehouse and BI solutions based on the BI Reference configurations include:

- Simplified product selection and sizing process
- Pre-defined set of core components to simplify purchasing and identify volume purchase opportunities
- Simplified, rapid deployment with predictable performance
- Off the shelf components to enable re-purposing and lower inventory risk
- Proven, award-winning support, with access to performance and tuning services

In summary, HP Reference Configurations for Oracle 10g Data Warehousing provide optimal BI performance and enable customers to avoid common design pitfalls.

HP advantage

As reported by IDC, HP is the world leader in infrastructure for strategic BI, and Oracle is the leading DBMS for data warehousing.¹ HP has been named as the best provider of BI servers.² Together, HP and Oracle support many of the world's largest data warehouse deployments, as well as tens of thousands of smaller implementations. With HP BI Reference Configurations for Oracle 10g Data Warehousing, customers can quickly and easily deploy high-performance business intelligence solutions for accelerated time-to-intelligence, reduced costs, and increased revenue.

¹ Server Workloads 2005: Understanding Server Deployment – IDC, 6/05 & Storage Workloads 2005: Understanding Storage Deployment – IDC, 9/05.

² 2005 Readership Award Winners, DM Review, December 2005.

For more information

HP Business Intelligence Solutions

www.hp.com/go/bi

HP and Oracle Grid Accelerator Service

www.oracle.com/technology/consulting/10gservices/hp.htm

Next steps

Putting HP BI Reference Configurations to work for you

Let HP help you get started. Contact your HP sales representative for more information on how a solution based on BI Reference Configurations can improve time to deployment for Oracle-based data warehousing. Your sales representative can also arrange for a TCO assessment to compare the performance of a BI Reference Configuration relative to your existing environment.

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

4AA0-6790ENW, 07/2006

